

The LANSCE Proton Storage Ring Sets a New Record	64
Unique Long-Life Carbon Stripper Foils Used in the Proton Storage Ring	66
Novel MOSFET Shunt Design Provides High-Precision Current Sharing	68
Continuous-Wave Radio-Frequency Quadrupole for the Low-Energy Demonstration Accelerator Successfully Reaches Design Beam Power	70
First Experiment to Study Halo in a Proton Beam at the Low-Energy Demonstration Accelerator	74
Collector Failures on 350-MHz, 1.2-MW Continuous-Wave Klystrons at the Low-Energy Demonstration Accelerator	78
Operational Results of the Spallation Neutron Source Polyphase Converter-Modulator for the 140-kV Klystron Radio-Frequency System	82
New Tools to Transport and Archive Data in EPICS	86
AccelView—An Innovative Tool Designed to Simulate, Visualize, and Animate Accelerator Beam Lines	90
H ⁻ Ion Source Development at LANSCE	94
Chopper Control	96

The LANSCE Proton Storage Ring Sets a New Record

R.J. Macek (LANSCE Division)

In November 1999, LANSCE achieved record-breaking levels of stored charge in its PSR. During the November 1999 tests, 9.7 μC of charge was successfully accumulated and stored in the PSR. This amount of charge corresponds to a circulating peak current of 82 A and represents nearly a 60% improvement over previous levels. The November tests were conducted as a part of the LANSCE Short Pulse Spallation Source (SPSS) Enhancement Project. SPSS project goals include a doubling of the available neutron flux at the Lujan Center and the construction of three new neutron-scattering spectrometers.

Upgrades Will Double Available Neutrons


The PSR—a circular ring 90 m in circumference—accumulates 800-MeV protons from the LANSCE linac and then delivers them in a short, intense burst to a tungsten target at the Lujan Center. This proton pulse produces spallation neutrons in the target. These neutrons are used by researchers at the Lujan Center for a wide variety of defense and civilian applications. To enable high-precision neutron-scattering measurements, the PSR compresses up to 1 ms worth of accelerator protons into just a quarter of a microsecond, providing a pulse 3000 times as intense in just one three-thousandth of the time.

The number of neutrons available to researchers at the Lujan Center is directly proportional to the number of protons delivered to the spallation target. The SPSS Enhancement Project will double the neutrons available by upgrading the PSR and the linear accelerator's H⁻ ion source. The PSR presently provides an average current of 100 μA at 20 Hz to the Lujan Center target, which corresponds to 5- μC pulses delivered 20 times per second. The upgrade will raise these pulses to 6.7 μC , which when delivered at a 30-Hz rate will provide an average current of 200 μA to the target. The ion source upgrade will reliably provide more protons to the PSR, and the PSR upgrade will enable the ring to handle the increased charge while maintaining beam losses at an acceptable level.

During the summer and fall of 1999, the SPSS Enhancement Project conducted a series of tests to demonstrate the PSR's ability to accumulate and store circulating beam at levels higher than ever before. Since

its commissioning in 1988, the upgraded PSR charge levels have been limited by a fast transverse-beam instability. At higher charge levels this PSR instability, if uncontrolled, causes sudden loss of all circulating beam. The aim of the 1999 tests was to demonstrate that the instability could be controlled at SPSS levels of 6.7 μC and above using a combination of the recently upgraded radio-frequency (rf) buncher, two sextupole magnets, and two inductive inserts. The tests are also expected to validate and extend the results of previous tests conducted in 1997 and 1998 at levels below 6 μC .

To date, the project team has demonstrated consistent control of the PSR instability at levels above 8 μC (Fig. 1). The highest test levels have reached 9.7 μC per pulse with beam losses around 4% to 5%. The previous record, set in 1988, was 6.1 μC with unacceptably high losses of over 6%. Low beam losses are required to minimize radiological activation of PSR components. Component activation poses hazards to service personnel. Future high-intensity tests will include special measures to reduce beam losses.



↑ Fig. 1. From the controls of the PSR, LANSCE scientists (l to r), Bob Macek and Dan Fitzgerald, review some test data, which reveal that the PSR has reached 8.7 μC , a significant step on the road to the record of 9.7 μC per pulse.

In the course of achieving the record intensity, the project team encountered an undesirable longitudinal resonance introduced by the inductive inserts. At the suggestion of Milorad Popovic of FERMI National Acceleration Laboratory (FERMILAB), the ferrite was heated, which successfully eliminated the longitudinal resonance. The engineered version of the inductive inserts (Fig. 2) includes heating elements that hold the temperature of ferrite of the inductors at the right level to eliminate the resonance.

The PSR high charge accumulation tests were conducted by the PSR Development Team, led by Bob Macek of LANSCE Division, with assistance from the LANSCE Operations Team. Collaborators who were instrumental in planning for the tests and providing hardware were Jim Griffin and David Wildman of FERMILAB, along with Bob Kustom and Kathy Harkay of ANL.



↑ **Fig. 2.** Photo of the engineered version of the two-module heated inductive insert permanently installed in Section 5 of the PSR. The inductive insert compensates the longitudinal space charge forces and thereby keeps beam from leaking into the gap in the beam bunch. This helps raise the instability threshold and also allows stable accumulation of a longer bunch in the ring. The later reduces the accumulation time, which saves electric power and also helps reduce beam losses in the ring.

Research Highlights

Unique Long-Life Carbon Stripper Foils Used in the Proton Storage Ring

M.J. Borden (LANSCE Division), I. Sugai (High-Energy Accelerator Research Organization, KEK)

Charge-exchange injection of beam into the PSR is accomplished by charge stripping in a thin carbon foil. Carbon foils produced with a modified and controlled ac and dc (alternating current/direct current) arc-discharge (mCADAD) method have significantly improved the operational performance of the PSR. The decreased shrinkage rates associated with these foils allow the use of smaller width foils, thereby decreasing stored beam losses. Lower beam losses imply that the ring components are less radioactivated, and therefore worker radiation exposures are reduced. Initial testing of foils produced by the mCADAD method, before the LANSCE Reliability Improvement Project, demonstrated lifetimes an order of magnitude higher than commercially available foils.

Unique Evaporation System

Over the past four years, personnel in LANSCE-2, under the direction of Michael Borden and the guidance of Dr. Sugai, have assembled a carbon-stripper-foil evaporation system that can produce foils with the mCADAD method. The system consists of a large stainless-steel vacuum chamber, cryo-pump, rotating substrate holder,



↑ Fig. 2. Electrode holder (center); annealing furnace (right).

arc-welding equipment, and annealing power supplies (Fig. 1). The electrode holder and annealing furnace may be seen in Fig. 2. The carbon electrode holders are machined from Macor™. Thickness is measured *in situ* during deposition with an oscillating quartz crystal monitor. Annealing temperatures are measured with type-K thermocouples placed in contact with the substrate glass and copper homogenizer.

Carbon stripper foils produced by the mCADAD have been tested and used for high-current, 800-MeV beam production in the PSR since 1993. The mCADAD foils are produced by arc-evaporating alternating layers of carbon produced with ac and dc. The carbon is deposited on clean glass substrates and then annealed to remove residual stresses. Foils in the mass density range of 100 to 130 $\mu\text{g}/\text{cm}^2$ each are sandwiched together to produce an equivalent 400- $\mu\text{g}/\text{cm}^2$ foil. The foil sandwich is supported by 4- to 5- μm -diam carbon fibers attached to an aluminum frame.

Foils produced by the mCADAD method in the 200- $\mu\text{g}/\text{cm}^2$ range have survived as long as six months during PSR pre-improvement beam production of near 70- μA average current on target. Typical lifetimes of other foils vary from seven to



↑ Fig. 1. mCADAD fabrication system.

fourteen days with lower on-target average current. The mCADAD technique is limited only in its ability to produce foils with a mass density in the range 50 to 130 $\mu\text{g}/\text{cm}^2$. Work is currently in progress to improve on this limit. With this limitation, a minimum of three foils is necessary to achieve an equivalent thickness of 400 $\mu\text{g}/\text{cm}^2$ for operation at 100- μA average current.

Both three-layer and four-layer $\sim 400\text{-}\mu\text{g}/\text{cm}^2$ assemblies have been used during normal 100- μA beam production. Average foil lifetime during the 1999–2000 beam-delivery period was approximately twenty-seven days, which is an acceptable lifetime for beam operations and continues to be optimized.

Novel MOSFET Shunt Design Provides High-Precision Current Sharing

S. Cohen, L.L. Duncan, J.E. Griego, E.A. Morgan, M. Fresquez, A.D. Archuleta (LANSCE Division)


Metal-Oxide Semiconductor-Field-Effect Transistor (MOSFET) circuit components are used routinely in industrial applications to control, for example, electric trains, induction furnaces, high-power radio transmitters, and other industrial equipment. MOSFETs can serve as shunts, which are electronic variable resistors used to regulate current in a power supply. Electric currents, hundreds to thousands of amperes, can be controlled using multiple MOSFET circuit components in parallel. Ensuring that currents are shared equally among multiple MOSFETs requires advanced, precise technology. We have developed a novel closed-loop circuit design that senses and adjusts current so that it is shared equally between parallel MOSFETs. Equipment that employs these components is used in LANSCE accelerator operations and run continuously for months at a time. The mechanical design and self-diagnostics of the closed-loop system have been engineered so that the average time it takes for components to fail, known as the mean time between failures, is very long (high), and the average time it takes to repair failed components, known as the mean time to repair, is very short (low).

Power-Supply-Regulator Passbanks and the Proton Radiography Shunts


A number of magnets in the LANSCE beam line require power supplies that provide a constant current with regulated stability of up to ten ppm (0.001%) to the magnets. Once accelerator operators and physicists set a specific current in a magnet, the power supply must maintain current within a tolerance of 0.001%. Devices with this strict performance specification use a passbank or electronic load resistor to provide the required precise current regulation. The passbank consists of many transistors or MOSFETs connected in parallel. The challenge of using multiple circuit components is to have them share current equally.

Currents on the order of hundreds or thousands of amperes need to share current equally because each individual MOSFET cannot carry the total current by itself. One way to ensure that current is shared equally is to use semiconductor devices with performance characteristics that are nearly identical. However, this approach is very expensive because it requires a large inventory of semiconductor devices.

The LANSCE-6 Magnet Power Supply Team took an alternative approach to facilitate equal current sharing among the MOSFETs by designing each component with separate closed-loop control circuitry (Figs. 1 and 2). This paradigm allowed us to use more generic parts because a feedback loop on each device guarantees even current distribution. The design includes features that allow easy maintenance, testing, and repair.




↑ Fig. 1. An operator sets the current required on the control panel. An electrical reference signal is sent to each of the closed-loop control circuits, which then adjust the MOSFETs electrical resistance to allow just the right amount of current to flow.



↑ Fig. 2. Schematic diagram of the placement of the MOSFET shunt in parallel with the quadrupole magnet, which is connected in series.

The entire assembly is modular in design so that a single inoperative MOSFET can be replaced without the need to disassemble nearby components (Figs. 3 and 4). Testing can be done safely without any disassembly or removal of the passbank from its mounting.




↑ Fig. 3. MOSFET passbank with closed-loop current-sharing circuitry. The MOSFET components are mounted on a water-cooled aluminum rail. The modular design and ribbon cable allow an arbitrary number of components to be wired in parallel to meet a variety of current-handling requirements.

Brookhaven E933 Magnet Shunt System

The MOSFET passbank design was adapted for use as a remote-control dual-channel current shunt for the PRAD experiment at BNL.

Because the shunt operated about 100 V above ground, for safety reasons the remote-control circuitry is electrically isolated from the shunt assembly. The shunts can handle 100 A per channel. The entire system was built into a single 19-in. rack and was shipped to BNL for installation, adjustment, and testing. The shunts were used successfully by the LANL PRAD collaboration during the August 1999 E-933

experimental run. The MOSFET passbank design scheme is being used successfully for 100-A shunts to control the quadrupole magnets in the 201-MHz drift-tube linac at LANSCE.



↑ Fig. 4. MOSFET passbank close-up. Note that each feedback circuit on the printed circuit board is mounted directly above the MOSFET components. The tinned wire loop (right center) is a test point for a non-contacting clamp meter.

Continuous-Wave Radio-Frequency Quadrupole for the Low-Energy Demonstration Accelerator Successfully Reaches Design Beam Power


L.M. Young, A.H. Regan (SNS Division), L.J. Rybarczyk, D.L. Schrage, J.D. Sherman, D.E. Rees (LANSCE Division), J.D. Schneider, H.V. Smith, Jr. (AAA-TPO), M.E. Schulze (General Atomics Corporation)

For the Low-Energy Demonstration Accelerator (LEDA),¹ LANSCE supported the design and construction of a 350-MHz continuous wave (CW) radio-frequency quadrupole (RFQ).² This accelerator successfully accelerated, for the first time ever, a 100-mA CW proton beam from 75 keV to 6.7 MeV for more than eight hours. To prepare the RFQ to accelerate the proton beam, we used a process known as high-power rf conditioning to clean and condition the surfaces in the RFQ. This technique is important in reaching high peak fields in most types of rf accelerators. When operating the RFQ at the rf power level for which it was designed, the peak electrical field on the vane tips is 33 MV/m. The RFQ dissipates 1.2 MW of rf power. At this power level, the average rf power dissipation is 13 W/cm² on the outer walls near the high-energy end of the RFQ.

We conditioned the RFQ with the high-power rf at power levels in excess of 1.4-MW CW and with pulsed beam before we obtained reliable operation with the 100-mA CW proton beam rf power from a low-level rf (LLRF) system is amplified by three klystrons. The amplified rf power then travels to the RFQ through waveguides. Each klystron has a very large power rating of 1.3 MW. To prevent damage to the waveguide vacuum windows, we split the power from each klystron two ways. These vacuum windows isolate the air in the waveguide from the vacuum in the RFQ, but they allow the rf power to pass through.

LEDA Configuration

LEDA was constructed to confirm the design and demonstrate the viability (at full power) of the most technically challenging components of the Accelerator Production of Tritium (APT) plant accelerator. On December 22, 1998, Secretary of Energy Bill Richardson announced that commercial light-water reactors will be the primary tritium-supply technology, and he designated APT as the backup technology for tritium supply. The LEDA RFQ consists of four 2-m-long RFQs resonantly coupled together to form an 8-m-long structure (Fig. 1). These segments are labeled A, B, C, and D starting from the low-energy end. The rf power is coupled into the




↑ Fig. 1. A solenoid magnet (not shown) focuses beam from the ion source into the low-energy end (Segment A) of the RFQ. As the beam travels through the RFQ gaining energy, the vane-tip modulation wavelength becomes longer. Radio-frequency power enters the RFQ at Section B1 (and D1) through small slots, known as irises, which are located between the ridge-loaded waveguides and the RFQ. The RF electric field accelerates and focuses the beam as it travels through the RFQ.

RFQ through a half-height WR2300 waveguide and a section of tapered, ridge-loaded waveguides connected to a coupling iris on the B1 and D1 segments. The tapered waveguide has the dimensions of the half-height WR2300 waveguide at one end and tapers to only 7 in. wide at the iris.

Iris Melting

On January 5, 1999, we achieved 1.2 MW of net rf power in the RFQ. However, the reflected power on the D-segment waveguide network increased to 10% instead of the 2% observed on December 24, 1998, the date of the previous test. Also, the frequency of the RFQ was 80 kHz lower than previously measured. While looking for the cause of the high-reflected power, we removed the window on quadrant 3 of segment D and discovered that some melting had occurred at both ends of the iris. Fig. 2 is a picture of an iris removed from the RFQ in April 1999. A MAFIA code calculation on a simplified model of the iris showed the rf currents at the ends of the iris slot were enhanced by a factor of 10 over the RFQ wall current. The copper iris plate was only 1/16 in. thick at this point. This localized heating coupled with the increase in surface resistance at elevated temperatures was enough to melt the ends of the iris slots when the RFQ was operated with CW rf power at close to the design fields. The model results also showed that by increasing the thickness of the iris plate to 3/8 in., the size of the hole at the end of the slot more than doubles for the same coupling. This change reduces the enhancement of the wall currents from a factor of 10 to only 2.5. We replaced the original 1/16-in.-thick iris plates with 3/8-in.-thick iris plates (Fig. 3).



↑ Fig. 2 Photo of melted iris.




↑ Fig. 3. Photo of new iris modified to reduce rf heating.

Pulsed Beam

After January 27, 1999, but before the iris plates were replaced, the average rf power was restricted to less than 1 MW to prevent further damage to the irises. This meant that the rf power had to be pulsed whenever we raised the rf fields to the design value or above. During this period we integrated the LLRF, the HPRF, and the water-cooling system to reliably maintain the rf fields in the RFQ. The LLRF system sends a frequency error signal to the RFQ's water-cooling system. This cooling system adjusts the temperature of the water that cools the outer walls of the RFQ to maintain the 350-MHz resonant frequency. The cooling system also supplies 50°F water to cooling channels near the vane tips. The first pulsed beam was obtained on March 16, 1999. The beam current was increased slowly until April 8, 1999, when the RFQ accelerated a 72-mA pulsed beam at 10 Hz. The pulse length was 1.5 ms. This current was the highest obtained with two klystrons driving the RFQ.

Continuous Wave Beam Tests

On June 11, 1999, after reconfiguring the system and after two weeks of rf conditioning, we operated the RFQ CW with the design fields for more than 30 minutes. The ion source was reconnected to the RFQ and testing with a pulsed beam resumed on June 18, 1999. We began CW beam tests August 9, 1999, and by September 17, 1999, we had accelerated 100 mA to 6.7 MeV CW. The maximum beam current that we were able to reach either pulsed or CW did not show any substantial improvement from August 10-24, 1999. Up to this time we were not able to make meaningful transmission measurements because the measured input current was less than the measured current out of the RFQ. Electrons flowing into the RFQ



↑ **Fig. 4.** RFQ output beam current versus time for a 300-μs-long pulse at ~ 97% of the design rf field level.

reduced the positive proton current measured by the toroid at the RFQ entrance. These electrons are responsible for neutralizing the space charge caused by the electrically charged protons. The presence of this space charge produces a significant electric field that tends to spread the beam. If electrons continue to flow into the RFQ, then the space charge is not being neutralized as well as it should be.

Using the computer code PARMELA, we performed a simulation of the beam traveling through the low-energy beam transport (LEBT) system with the proton beam's space charge neutralized at 98%, except for the last 10 cm in front of the RFQ. The results of this simulation showed that the beam could not be "matched" to the RFQ. To be matched, the beam must be the correct size and converge at the correct angle when it enters the RFQ. Space charge caused the beam to deflect so much that it no longer converged as it entered the RFQ. In this situation, the beam was "mismatched," which means that the beam was too small in some locations and too large in other locations where it was scrapped off on the vane tips. The simulation further showed that the maximum beam current out of the RFQ with this LEBT configuration was only 89 mA, which is equal to the maximum pulse current that we obtained by August 23, 1999. We moved the solenoid that focused the beam into the RFQ closer to the RFQ in an effort to match the beam into the RFQ. The solenoid was 30 cm from the RFQ. Simulations showed that if we installed an electron trap to prevent the electrons from flowing into the

RFQ and moved the solenoid 15-cm closer to the RFQ, then the beam could be matched into the RFQ. The electron trap is a ring with a negative 2-kV potential placed at the entrance of the RFQ through which the beam passes. The potential from this ring prevents electrons from going through it, but not protons. We made these changes to the LEBT on August 28-29, 1999. A water leak in the LEBT beam stop prevented us from running beam until September 7, 1999. But then we were able to reach our beam current goals in only ten days after this change. We also were able to obtain realistic measurements of the transmission through the RFQ. The electron trap stopped electrons flowing into the RFQ and stopped the electrons from reducing the measured input current.

An unexpected reduction in transmission, shown in Fig. 4, is observed in a 300-μs-long beam pulse when the RFQ fields are near the design field level. The transmission drops unexpectedly about 150 μs into the pulse. As we raise the RFQ fields the transmission stays high for longer times. With fields above ~105% of design we no longer observe this drop, even for long pulses and CW operation. We now operate the RFQ routinely with the RF fields about 10% above design field level to avoid this reduction in transmission.

Summary

On December 17, 1999, LEDA operated with CW beam current of ~100 mA with only a few short

interruptions and averaged 98.7 mA over 3.3 h.³ The LEDA RFQ performs as designed, provided the rf field is raised about 10% above the design level. The addition of an electron trap at the entrance to the RFQ is essential to the measurement of the transmission through the RFQ. Simulation of the beam transport through the LEBT with PARMELA allowed understanding the injection of the beam into the RFQ.

The high-energy beam transport and beam stop have been moved ~11 m to make room for a 52 quadrupole beam transport line. The RFQ is now operating with beam through the new beam line. This beam line is being used to study beam-halo of both matched and unmatched high-current beams.⁴

References

1. J. D. Schneider, "Operation of the Low-Energy Demonstration Accelerator: The Proton Injector for APT," in *Proceedings of the 1999 IEEE Particle Accelerator Conference* (IEEE, Catalog No. 99CH36366, 1999), 503-507.
2. D. Schrage *et al.*, "CW RFQ Fabrication and Engineering," in *Proceedings of the Fourteenth International Linac Conference* (Argonne National Laboratory, ANL-98/28, 1998), 679-683.
3. L. M. Young *et al.*, "Low-Energy Demonstration Accelerator (LEDA) Radio-Frequency Quadrupole (RFQ) Results," in *Proceedings of the ICFA Advanced Accelerator Workshop on the Physics of High-Brightness Beams* (Microsoft Press, Redmond, Washington, 1995).
4. T. P. Wangler, "Beam Halo in Proton Linac Beams," in *Proceedings of the Twentieth International Linac Conference*, Monterey, California, August 21-25, 2000.

First Experiment to Study Halo in a Proton Beam at the Low-Energy Demonstration Accelerator

R.W. Garnett, D. Gilpatrick, W. Lysenko, J. Qiang, T.P. Wangler, P. Colestock (LANSCCE Division), K.C.D. Chan, J.D. Schneider, R. Sheffield, H.V. Smith (AAA-TPO), C.K. Allen (SNS Division), M.E. Schulze (General Atomics Corporation), K.R. Crandall (TechSource)

Beam halo is an important characteristic of high-intensity beams. Under certain conditions in proton beams, some of the particles can acquire enough transverse energy from the repulsive space-charge forces within the beam to form a halo at the periphery. The formation of halo is an important concern because if the halo particles are lost at high energies on the walls of the beam-line structures, they will induce unwanted radioactivity. Because the radioactivation from the beam losses complicates the maintenance of the accelerator, beam halo can cause a significant reduction in the availability. Therefore it is crucial to gain understanding of the physical causes and possible cures for beam halo and beam losses. The next generation of proton-linac applications requires beams with higher average currents than the LANSCCE linac produces today. These higher average beam currents increase the risk of beam losses. Ultimately, we need an experiment that tests and demonstrates our capability to predict and suppress the beam halo.

Studying Beam Halo Formation

We are installing a beam line at the end of the LEDA¹ to carry out a first experimental study of beam halo formation in a proton beam. The LEDA facility will deliver a 6.7-MeV, 100-mA proton beam from the 350-MHz RFQ² to the new beam line. The beam line will consist of a periodic array of 52 quadrupole magnetic lenses for transverse confinement of the beam and a number of beam-diagnostic devices to monitor and measure the beam properties. The halo experiment will require measurements of the horizontal (x) and vertical (y) distributions of the beam particles.

To study the beam halo, we will set the quadrupole strengths to deliberately establish conditions that we believe are favorable for halo formation, and we will measure and analyze the resulting beam profiles. The measured halo effects will be compared with those predicted by our beam-dynamics simulation codes, and the comparisons will enable us either to confirm that our computer models are correct or will help us identify possible missing physics effects that may need to be incorporated. The measurements will answer practical questions concerning whether the apertures in high-current accelerator designs are properly chosen, and

whether or not special collimation systems are required at low energies to scrape the halo and limit the high-energy beam losses.

Objectives of the Halo Experiment

There are two main objectives of the experiment. The first objective is to test the predictions of the beam-dynamics simulation codes, particularly the space-charge subroutines, against the transverse beam-profile measurements. The LEDA beam-line parameters can be varied to establish different known conditions that are more or less favorable for formation of halo. We are interested in determining the maximum amplitudes of the beam particles so that we will know how to choose the beam apertures in future linac designs to avoid beam loss. We are also interested in determining how the particles are distributed throughout the beam cross section so that we know the relative population of the particles at different radii in the beam, including the outer halo. As a check on our understanding of the physics, we are also interested in measuring the emittance of the beam. The emittance is an important parameter that is determined by the area occupied by the beam particles in phase space, which is a space that is defined by the coordinates and angles of the particles relative to the beam center. The emittances can be determined experimentally from multiple measurements of the beam profiles at different locations.

Our second experimental objective is to explore those regions beyond where halo is predicted by simulations to ensure that there are no other important sources of halo than those predicted by the simulation codes. After the measurements are made and the data are analyzed, our goal is to be able to choose confidently the beam-line apertures for new accelerator designs, or, if necessary, to be able to design beam-collimation systems to remove the halo so that beam losses and radioactivation are small effects.


We believe the main cause of beam halo in a periodic focusing channel is beam mismatch. A beam is matched when there is a balance between the inwardly directed focusing force and the outwardly directed forces composed of the space-charge force plus a pressure force from the emittance. Beam mismatch

(not shown) in which the beam is focused by four quadrupole magnets and transported to a beam stop. The output beam-transport system has additional beam-diagnostic components to ensure that the beam is properly directed onto the beam stop.

In the layout of Fig. 1, beam-profile scanners are distributed in two clusters of four scanners each. The choices of four scanners per cluster, the spacing of the scanners, and the locations of the clusters allow us to calculate the emittances from the beam-profile measurements and obtain both upstream and downstream values. The first cluster of four is located shortly after the beam has debunched, and the second cluster is located near the end of the beam line. The scanners are spaced apart by one focusing period (two quadrupoles). The first scanner is installed near the beginning of the beam line to monitor the input particle distribution and to determine whether there is halo present in the input beam coming from the RFQ.

Beam Profile Scanners

With its unique design, the beam-profile scanner (see Fig. 2) is the most important diagnostic device for the experiment. Each scanner contains a single 33- μm -diam graphite wire for measurement in each transverse plane of the high-intensity central beam core. In addition, each scanner contains a pair of copper plates, one on each side of each wire to intercept and measure the beam intensity in the low-density halo region. Beam particles that pass through the wire knock out secondary electrons. A resulting current is induced in the wire, proportional to the number of beam particles that are intercepted by the wire. The induced current is measured by the electronics. Any beam particles that impinge on the plates are stopped in the plates, which are thicker than the particle range in the copper. The stopped beam induces a current in the plates that is also measured by the electronics. The use of the plates instead of a wire for measurements in the halo is to collect beam over a larger area, which increases the electronic signal-to-noise ratio. The peak beam flux delivered to any material in the pulsed beam exceeds 1 MW/cm², sufficient to cause melting and material damage. Therefore, it is intended that the plates are never to be moved into the intense core of the beam, where they would be damaged. The wire/plate assembly for each scanner in each plane is moveable by external actuators and is stepped through the beam to obtain a complete measurement of the beam profile. The wire/plate assemblies are fixed in position during a single macropulse while the measurement is being made. Generally, only one wire scanner at a time is positioned to intercept the beam. The



↑ Fig. 2. Drawings of the beam-profile scanners. The drawing on the top shows the scanner assemblies for both planes and the drawing on the bottom shows the assembly for the vertical scanner alone. Each scanner assembly contains a wire and a pair of plates (both are shown in red). Not all the connections that support the wires and plates are shown. The beam pipe size is shown by the circle. The assembly in each plane can be moved between macropulses along a 45° angle to reposition the wire and plates for each new measurement.

wire/plate assembly can be stepped to a new transverse position during the time interval between macropulses. A complete set of measurements yielding the complete beam profile in each transverse plane can be made within approximately one minute per plane per scanner. Each scanner will be aligned in both x and y to an accuracy of $\sigma = \pm 0.0254$ mm. The wire will provide intensity measurements over a dynamic range of approximately 1000. The plates will extend the lower limit of this range by another factor of 10 to 100 for the halo measurements. For a beam with a Gaussian profile, the wire provides a transverse profile measurement out to approximately four standard deviations, and the plate extends this to approximately five standard deviations.

Summary

The beam-halo experiment at LEDA will provide the first experimental tests of our understanding of halo formation in high-current proton beams. We have carried out multiparticle simulations that provide theoretical predictions. Nevertheless, past experience cautions us to recognize that, given the present state of the art

of multiparticle simulations and beam measurements, one does not expect exact agreement between the measurements and the simulations shown above. One issue is unavoidable approximations in the simulation codes, which may not be accurate enough. This is why we use more than one simulation code. Another issue is possible missing physics in the codes, such as charge neutralization from beam-induced ionization of the background gas. With respect to this charge-neutralization issue, we believe that the short pulse length and the high vacuum should ensure that this is unimportant. Perhaps the most important uncertainty is the lack of detailed knowledge in all three dimensions of the input beam from the RFQ, especially the longitudinal beam properties, which are very difficult to measure. These are all reasons why we need an experiment before we have confidence necessary for the design of proton linacs for future applications requiring very high average beam currents.

The 52-quadrupole beam line and the beam diagnostics are being built, tested, and installed. Experimental measurements will take place during 2001.

References


1. J. D. Schneider, "Operation of the Low-Energy Demonstration Accelerator: The Proton Injector for APT," in *Proceedings of the 1999 IEEE Particle Accelerator Conference* (IEEE, Catalog No. 99CH36366, 1999), 503-507.
2. D. Schrage *et al.*, "CW RFQ Fabrication and Engineering," in *Proceedings of the XIX International Linac Conference* (Argonne National Laboratory, ANL-98/28, 1998), 679-683.
3. T. P. Wangler, *Principles of RF Linear Accelerators* (John Wiley & Sons, 1998), pp. 270-272.
4. J. Qiang, R. D. Ryne, S. Habib, and V. Decyk, "An Object-Oriented Parallel Particle-in-Cell Code for Beam Dynamics Simulation in Linear Accelerator," *Journal of Computational Physics* **163** (2), 434-451 (2000).

Collector Failures on 350-MHz, 1.2-MW Continuous-Wave Klystrons at the Low-Energy Demonstration Accelerator

D. Rees, W. Roybal (LANSCE Division), J. Bradley III (SNS Division)

The Accelerator Production of Tritium (APT) 350-MHz, 1.2-MW CW klystrons operating on LEDA are similar in design to the English Electric Valve (EEV) 352-MHz klystrons at the European Center for Nuclear Research (CERN), the Advanced Photon Source (APS), and the European Synchrotron Radiation Facility (ESRF) with one difference. The APT klystrons were designed to allow for the steady-state dissipation of the full direct-current (dc) beam power (95 kV, 20 A, 1.9 MW) in the collector. This requirement was intended to mitigate alternating-current (ac) grid transients that could result from loss of accelerator beam since the APT accelerator is heavily beam loaded or from interlock-induced, short-term accelerator outages. Each klystron was factory tested to this requirement and operated at LEDA without an interlock to limit the time the collector was subject to the full beam power in the event the rf drive to the klystron is interrupted.

Three klystrons of this type provide power to an RFQ. Approximately 2.4 MW of rf power is required from the klystrons for the RFQ. Three klystrons are connected to the RFQ, which acts as the power combiner. The system is designed so that only two of the three klystrons are required for operation. Waveguide switches are included in each waveguide run so that a failed klystron can be removed. The switch reflects a short circuit at the appropriate phase back to the RFQ and the other two klystrons can be used to continue operations. A picture of the klystron is included in Fig. 1.



↑ Fig. 1. EEV 350-MHz APT klystron.

Discovery of Damage


Initial Indications. The first indication of a problem was an audible boiling in the collector. The collector is a mixed-phase collector that allows for boiling and

recondensation within the collector grooves. The audible boiling was not this mixed-phase boiling. The noise was the result of large steam bubbles being transported at high water velocities through the collector cooling pipes and metering. We notified the vendor of the boiling and were told it was nothing to worry about. This boiling was noticeably worse on one of the three klystrons but was observable on all klystrons for high collector powers.

In order to keep the audible boil to a minimum, we had been operating the klystrons at the minimum beam power required for RFQ operations. Earlier in the day we had increased the klystron beam power from approximately 1.1 MW to 1.5 MW to allow for increases in RFQ beam current and provide more margin for field control. We had operated several hours at this new operating point with negligible ion-pump current (less than 2 μ A) when we experienced a large vacuum event that triggered both the high-voltage interlock (which triggers at 10 μ A ion-pump current) and the filament power interlock (which triggers at 100 μ A of ion-pump current). We suspended operations and attempted to recondition the klystron. Although we were able to recondition the klystron, as we operated at higher beam power, we began to observe variable cathode emission.

Variable Emission. As we reconditioned the klystron, we discovered that as we increased beam power, cathode current would vary by several amps when increasing or decreasing the rf drive. We could also cause this same variation by small changes in cathode voltage. Our system operates with a mod-anode regulator tube that maintains the mod-anode voltage at a fixed set point relative to the cathode voltage, and we carefully monitored this voltage, as well as the filament current, to verify neither of these factors were causing the change in cathode current. Finally, we did a series of tests where we triggered our arc detectors, which inhibit rf drive to the klystrons.

Fig. 2 shows the results of one of these tests. In the figure, we see the rf power go to zero and a corresponding plunge in cathode current. When the rf power is reestablished, we see the cathode current slowly return to its previous value. We can also see the filament current and mod-anode voltage relative to cathode voltage to be constant over the variation in cathode current. From this data, we concluded that copper vapor



↑ **Fig. 2.** Strip-chart record illustrating variable emission with changes in collector power.

from the collector was poisoning the cathode emission as we increased the thermal load in the collector. It was also interesting that the vacuum ion-pump power supply did not register an increase in vacuum and indicated very low current throughout this test.


Damage Assessment. Based on this result and conversations over the duration of our testing, EEV had two klystron engineers visit our site. We removed the end of the collector water jacket from the klystron and observed the damage shown in Fig. 3. From Fig. 3 we can see that several portions of the collector had become quite hot and dimples had resulted from a combination of the heat, water pressure, and vacuum forces. The general collector shape had also distorted. It was now elliptical in cross section rather than round. Also, the collector end had drooped, and the bottom of the conical shaped portion of the collector was actually resting on the water jacket.

Short-Term Solution

We inspected the collectors of all the klystrons, and all showed various levels of damage, ranging from minimal

to the significant damage shown in Fig. 3. All collectors showed indications of droop, and three of the four collectors were becoming elliptical rather than round in cross section. All collectors had less than 1000 high-voltage hours and one had only 60 high-voltage hours since delivery. Of the four klystrons, one klystron had a bucking coil in place to minimize magnetic-field penetration into the collector. This klystron had the least damage. Two other klystrons had this bucking coil installed but not electrically connected to the magnetic circuit. The fourth klystron did not have the bucking coil installed. We were dismayed to discover the variation in the magnetic circuit of the klystrons but were encouraged by the fact that the klystron with the operating bucking coil showed the least damage. We took the following short-term steps to return the klystrons to service due to programmatic pressure.

The klystron in Fig. 3 was returned to the factory for rebuild while still under vacuum and operable at low beam powers. The collector jackets of the remaining three klystrons were adjusted to center the collectors in the water jackets in order to compensate for the droop, and supports were welded on the inside of the water jacket to support the end of the collector to avoid further droop. An acoustical measurement system was installed on the collector to validate the effect of the bucking coil and to characterize safe



↑ **Fig. 3.** External view of collector damage.

Research Highlights

results. The mixed-phase boiling that the collector was designed to accommodate manifests as an increase in the higher frequency audio components of Fig. 4. This increase can be used as a qualitative indication of peak collector heat load. From Fig. 4 it can be seen that the localized collector heat load at 77 kV, 14.5 A is actually higher than at 85 kV, 17.1 A. This suggests that the beam power is better distributed over the collector at 85 kV. The plots in Fig. 4 were measured without rf power with the full beam power dissipated in the collector. Fig. 4 also shows that application of the bucking coil lowers the localized collector heat load. Using this tool, we were able to define an allowable safe operating region and also regions to avoid. Operation with three klystrons at an 85-kV voltage provides sufficient power for the RFQ with full accelerator beam current and sufficient control margin.


The data in Fig. 4 were measured on the klystron that had the second most significant amount of collector damage. The other two klystrons showed a similar

behavior but not the same magnitude of acoustical variation. For all klystrons, operation between 70 kV and 80 kV resulted in audio indications of increased localized heating.

In addition to the steps taken above, we also implemented a beam interlock such that above 70 kV of beam voltage, 100 kW of rf power was required or an interlock would remove high voltage within 15 s. We would have preferred a shorter interlock period, but the 15-s interval was a compromise with the RFQ operations team and with EEV.

Long-Term Solutions

EEV has redesigned the collector so it is capable of achieving the initial requirements. The details provided below are intentionally vague to protect proprietary information. Additional modeling was conducted using a recently released new version of one of the beam dynamics codes originally used in the collector design.



↑ Fig. 4. Collector audio profile.

The beam was discovered to penetrate further into the collector than originally thought. Additional length has been added to the collector in the new collector design. The conical collector end has been changed to a shape that reduces the thermal stress. The bucking coil is now a required part of the design. The water jacket design has been changed to maintain a high flow velocity over the entire collector. The supports on the end of the collector implemented as part of the short-term solution have become a part of the design. We also conducted x-ray measurements of the klystron internal to the lead shielding and discovered that the beam alignment could be improved. These results were duplicated by EEV, and the company is taking steps to improve beam alignment.

Conclusion

We discovered the initial collector design to be inadequate for satisfying our requirements. We have worked closely with EEV to diagnose and modify the collector design. The first 350-MHz klystron with the new collector design was scheduled for factory testing in

August 2000. The factory acceptance tests of the collector have been significantly increased, and an inspection of the collector is conducted at the conclusion of the testing. After delivery of this initial tube, another klystron from Los Alamos will be rebuilt to include the new collector. The remaining two klystrons at Los Alamos had only slight collector damage and will continue to be used on LEDA.

As part of diagnosing the collector problem, we developed a useful and interesting tool to characterize localized collector heat load using an acoustical measurement. This tool led to the unexpected result that lower beam powers had higher localized collector heat loads than operation at higher beam powers.

We have operated the three remaining klystrons for slightly over 1,000 cumulative hours since the collector problem was discovered. The two slightly damaged klystrons show no sign of further damage. The remaining klystron with the most significant damage showed signs of further degradation; however, it continued to meet the RFQ needs until the replacement tube arrived.

Operational Results of the Spallation Neutron Source Polyphase Converter-Modulator for the 140-kV Klystron Radio-Frequency System


W.A. Reass, P.J. Tallerico (SNS Division), J.D. Doss (NIS Division), R.F. Gribble (P Division), M.T. Lynch, D.E. Rees, D.L. Borovina (LANSCE Division)

This article describes the first operational results of the 140-kV, 1-MW average, 11-MW peak, zero-voltage-switching, 20-kHz polyphase bridge, boost converter-modulator for klystron pulse application. The dc-dc (direct current) converter derives the buss voltages from a standard 13.8-kV to 2100-V substation cast-core transformer. Energy storage and filtering is provided by self-clearing metallized hazy polypropylene traction capacitors. Three "H-Bridge" insulated gate bipolar transistor (IGBT) switching networks are used to generate the polyphase 20-kHz transformer primary-drive waveforms. The 20-kHz drive waveforms are chirped the appropriate duration to generate the desired klystron pulse width. Pulse-width modulation (PWM) of the individual 20-kHz pulses is used to provide regulated output waveforms with adaptive feedforward and feedback techniques. The boost-transformer design utilizes amorphous nanocrystalline material that provides the required low core loss at design flux levels and switching frequencies. Resonant shunt peaking is used on the transformer secondary to boost output voltage and resonate transformer leakage inductance. With the appropriate transformer-leakage inductance and peaking capacitance, zero voltage switching of the IGBTs is attained, minimizing switching losses. Reviews of these design parameters and an examination of the first operational results will be performed.

three-phase 13.8-kV to 2100-V vacuum-cast core transformer with passive harmonic traps and input harmonic chokes. These components are located in an outdoor-rated NEMA 3R enclosure that does not require secondary oil-containment or related fire-suppression equipment. The LANL prototype is manufactured by Dynapower Corporation in Burlington, Vermont. The substation is followed by a silicon-controlled rectifier (SCR) preregulator that accommodates incoming line-voltage variations and other voltage changes resulting from transformer and trap impedances, from no load to full load. The SCR preregulator also provides the soft-start function. The SCR preregulator provides a nominal ± 1250 -V output to the energy-storage capacitor banks. The SCR preregulator used in the LANL prototype was manufactured by NWL Transformers in Bordentown, New Jersey. The energy storage uses self-clearing capacitors developed for the traction industry. As in traction application, these capacitors are hard bussed in parallel. These capacitors do not fail short, but based on the use of metallized hazy polypropylene, they fuse or "clear" any internal anomaly. At our capacitor voltage rating (1.5 kV), there has not been a recorded internal capacitor buss failure. In this application, as in traction motor applications, the capacitor lifetime is calculated to be 400,000 h, including derating factors. A special low-inductance design for these capacitors has been developed by Thomson Components in Saint-Apollinaire, France. The IGBTs are configured into three "H" bridge circuits to generate a three-phase, 20-kHz, square-wave voltage-drive waveform applied to the transformer primaries. The IGBTs are "chirped" the appropriate duration to generate the high-voltage klystron pulse, typically 1.2 ms.

Component Descriptions

The simplified block diagram of the converter-modulator system is shown in Fig. 1. This system minimizes costs by using many standard and proven industrial and utility components. The substation is a standard



↑ Fig. 1. Simplified block diagram.

Because of the IGBT switching power levels, currents, and frequencies involved, low-inductance busswork and bypassing is of paramount importance. The IGBTs are 3300-V, 1200-A devices manufactured by EUPEC of Hanau, Germany. The boost transformers use amorphous nanocrystalline material that has very low loss at the applied frequency and flux swing. Operating at 20 kHz and about 1.6 T bidirectional, the core loss is about 1.2 W per pound in our application, or 320 W per core. Each of the "C" cores (one for each phase) weigh 260 lb and has a 3.5-in. by 5-in. post. The output high-voltage rectification circuit is a standard six-pulse rectification circuit with a "pi-R" type filter network. The diodes are high-voltage, fast-recovery ion-implanted types, manufactured by IXYS, which are series connected with the appropriate compensation networks. The diodes have the second highest total power loss (after the IGBTs) and are oil cooled.

Special Design Features of the Transformer and Insulated Gate Bipolar Transistor

By appropriately spacing the secondary from the primary, the transformer-leakage inductance can be resonated with secondary shunt-peaking capacitors to maximize voltage output and tune the IGBT switch current to provide "zero-voltage switching" with IGBT turn-on. The zero-voltage switching occurs when the IGBT gate drive is positive, but reverse transformer primary circulating current is being carried by the IGBT internal freewheel diode. We have tuned for about 2.5 μ s of freewheel current before the IGBT conducts in the normal quadrant. This tuning provides for about 10% control range (2.5/25 μ s) for IGBT PWM. Further transformer-design optimizations can change IGBT commutation (turn-off) current for control range and coupling coefficient for IGBT peak current. As transformer-design characteristics interact with other circuit parameters, optimization may be performed for various klystron loads and voltages.

Output Level Control

IGBT PWM of the active klystron voltage pulse enables us to use feedback and adaptive feedforward techniques with digital signal processors (DSPs) to regulate and provide "flat" klystron voltage pulses, irrespective of capacitor bank "start" voltage and related droop. The DSP adaptive feedforward/feedback processor used in the LANL prototype was manufactured by Z-TEC, Inc., of Albuquerque, New Mexico.

Fault Protection

The filter network must attenuate the 120-kHz switching ripple and have minimal stored energy. The stored energy is wasted energy that must be dissipated by the klystron at the end of each pulse. With the parameters we have chosen, the ripple is very low (~ 150 V), and the klystron fault energy (in an arc-down) is approximately 10 J. Even if the IGBTs are not turned off, the transformer resonance is out of tune in a fault condition, and little difference in klystron fault energy will result. If the IGBTs fail short through the transformer primary winding, the boost transformer will saturate in about 30 μ s, also limiting any destructive faults to the klystron. In a faulted condition, the klystron peak fault current is about twice nominal, with low di/dt s.

Modeling


The complete electrical system of the converter/modulator system has been modeled in extreme detail. This includes design studies of the utility characteristics, transformer and rectification methodology (e.g., 6-pulse vs 12-pulse), IGBT switching losses, boost-transformer parameters, failure modes, fault analysis, and system efficiencies. Various codes such as SCEPTRE, MicrocapIV, Flux2D, and Flux3D have been used to perform these tasks. SCEPTRE has been primarily used to examine IGBT and boost-transformer performance in great detail to understand design parameters such as switching losses, IGBT commutation di/dt , buss inductance, buss transients, core flux, core flux offset, and transformer Eigen frequencies. Flux2D and Flux3D have been used to examine transformer coupling coefficients, leakage inductance, core internal and external flux lines, winding electric field potentials, and winding field stresses. The Flux2D and Flux3D were particularly useful to examine transformer secondary winding profiles that gave the desired coupling coefficients with minimized electrical field stresses. MicrocapIV has been used to examine overall design performance of the system. This includes the utility grid parameters such as power factor, line harmonics, and flicker. We have optimized the design to accommodate the IEEE-519 and IEEE-141 harmonic content and flicker standards. Comparisons between the SCEPTRE and MicrocapIV codes show no significant differences in the system operational performance such as switching currents, switching voltages, and output voltage.

First Results


The converter-modulator made its first high-voltage pulse on January 17, 2001. After 10 days of testing and

Research Highlights

replacement of defective dummy-load components, full-pulse output voltage (140 kV), pulse width (~ 1.2 ms), and peak power (11 MW) were obtained. As shown in Fig. 2, the 140-kV risetime is about 20 μ s with about a 6% bank droop. The risetime compares favorably to other long-pulse modulator topologies. Additional high-voltage tests were made at 80 kV, as this voltage will be used in the SNS linac superconducting portion. The 80-kV operations are shown in Fig. 3.




↑ Fig. 2. The 140-kV output pulse, 20 kV/division.




↑ Fig. 3. The 80-kV output pulse, 20 kV/division.

Operations at 80 kV with the DSP adaptive feedforward/feedback processor is shown in Fig. 4. The DSP processor removes all overshoot and bank voltage droop. Extensive fault testing has also been performed with three times the anticipated high-voltage cable length as will be used in SNS.


The most severe crowbar test is the event of when the klystron arcs, but the converter-modulator continues switching the duration of the fault, causing a "run-on" fault situation. Modeling of this result shows a slight increase in klystron fault energy (to about 10 J) and a slight increase in IGBT current. The results of these



↑ Fig. 4. The 80-kV output with adaptive feedforward/feedback.



↑ Fig. 5. The 130-kV self-break crowbar test.



↑ Fig. 6. The 130-kV run-on fault test.

tests are shown in Figs. 5 and 6. Fig. 5 shows a 130-kV crowbar test with the IGBT switch drive immediately terminated at the fault event. Fig. 6 shows the same test, but the IGBT switching is not inhibited. The surprising result was that the IGBT over-current monitors

additionally did not trip. The low primary IGBT driving voltage cannot significantly increase primary IGBT current because the high transformer winding reactance cannot drive higher currents from the shorted and now resonant-tuned secondary winding. These results match our computer modeling very accurately.

Project Status

Low average power testing has been completely successful. The substation construction work is completed, and high average power testing began in May. We are presently testing on a 250-kW, 402-MHz klystron. Continued testing to the advertised 1-MW average power is of high importance. Testing at these power levels will commence once facility upgrades (such as cooling water) and loads (klystrons) become available.

Conclusion

The converter-modulator has demonstrated several new design methodologies that are expected to revolutionize long-pulse klystron modulator design. These items include special low-inductance self-clearing

capacitors, large amorphous nanocrystalline cut-core transformers, high-voltage and high-power polyphase quasi-resonant dc-dc conversion, and adaptive power supply control techniques. The first test results on the initial design were achieved in about a year after conception. Design economies are achieved by the use of industrial traction motor components (IGBTs and self-clearing capacitors) and standard-utility cast-core power transformers. The compact and modular design minimizes on-site construction, and a simplified utility interconnection scheme further reduces installation costs. The design does not require capacitor rooms and related crowbars. By generating high voltage only when needed, reliability and personnel safety is greatly enhanced. This approach provides design flexibility to operate klystrons of different voltages primarily by changing the boost-transformer turns ratio. Other optimizations also permit continuous operation of the polyphase boost converter topology. All testing of the full-scale system have been completely successful, and all results agree with modeling efforts to date, which indicate the design methodologies will be imminently successful.

New Tools to Transport and Archive Data in EPICS

K.-U. Kasemir (LANSCE Division)

We developed two new additions to the Experimental Physics and Industrial Control System (EPICS) toolkit. Originally deployed on LEDA,¹ they are now also available to other EPICS users. The first addition, an ActiveX ChannelAccess server, allows integration of Win32-based standalone subsystems, e.g., built-in LabVIEW, into distributed and portable EPICS environment. The second one, called ChannelArchiver, constitutes a high-performance archiving tool set for the EPICS network data protocol.

EPICS

EPICS is a set of reusable software tools for building distributed control systems. It was originally developed at LANL (for the ground-test accelerator project) together with ANL (for the Advanced Photon Source). Today more than 100 sites are part of the worldwide EPICS user base. Active collaborators include not only particle accelerators, but also large physics experiments and telescopes such as the SNS, LANL, Jefferson Laboratory, ANL, BESSY Laboratory, Stanford Linear Accelerator Center, the Keck telescope, DESY Laboratory, and others.

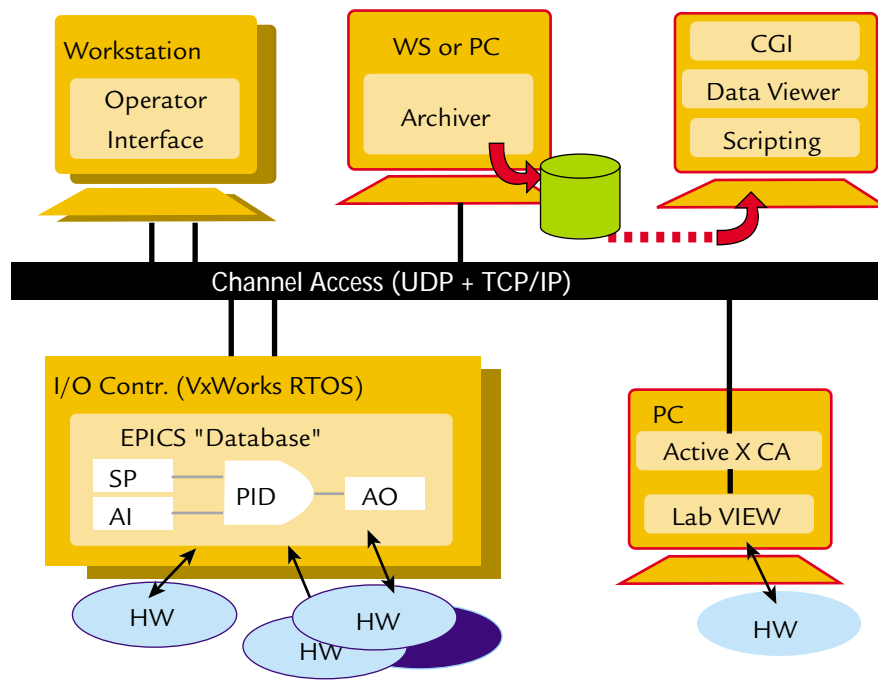
ChannelAccess (CA) is a publish-subscribe data protocol based on the user datagram protocol (UDP) and transmission control protocol/internet protocol (TCP/IP). It is the backbone of the distributed EPICS environment.^{2,3} Traditionally, front-end Input/Output (I/O) controllers (IOCs), based on the Versa Module Europe (VME) bus system and running the vxWorks real-time operating system, interact with the hardware and perform closed-loop control, configured via the EPICS "IOC Database." Over CA, they exchange data with each other, as well as with Unix workstations that run the graphical operator interfaces, alarm handlers, control scripts, and other programs as illustrated in Fig. 1.

For the LEDA project, some subsystems (beam-stop hydrophones, beam-position monitors, phase/energy measurements) were built as stand-alone LabVIEW applications and necessitated integration into the distributed environment. In addition, an archiving facility was required.

ActiveX ChannelAccess Server

The CA server C++ application program interface is available for Unix, as well as for Win32. For LabVIEW users on PCs running Microsoft Windows98, NT, or 2000, this could have been wrapped in a DLL. While this provides access to the complete CA functionality, it would put the burden of scheduling network traffic and transforming data types onto the LabVIEW programmers. Instead, the CA server library was wrapped into a standalone ActiveX Automation Server executable.

LabVIEW, as well as many other Win32 applications (most importantly Visual Basic), can access ActiveX Servers transparently through the Win32 Component Object Model (COM). As soon as one such application



↑ Fig. 1. Distributed components of EPICS; new components marked red.

tries to create an `EpicsCASServer.ProcessVariable`, the Win32 operating system launches the ActiveX automation server. The principal limitation of all TCP servers of allowing only one server per machine to use a well-known TCP port is overcome by the fact that multiple applications, e.g., several LabVIEW applications in combination with Visual Basic, can share the same ActiveX automation server.

Properties of the process variable include the name, value, engineering units, display precision, and limits. The value can be set as a double, long, or generic COM variant, including array types. The Automation Server time-stamps each value and serves all properties via CA to other CA clients on the network (for example, operator screens in the control room). When an operator changes the value on the display, the ActiveX CA Server transforms the incoming CA value into an ActiveX event for the Win32 application that created the process variable.

Several new custom virtual instruments (VIs) for LabVIEW further simplify the usage of the ActiveX CA server by presenting it similar to other LabVIEW instruments with init/set/check/close VIs (see Fig. 2). In addition, these VIs provide online help to LabVIEW programmers.

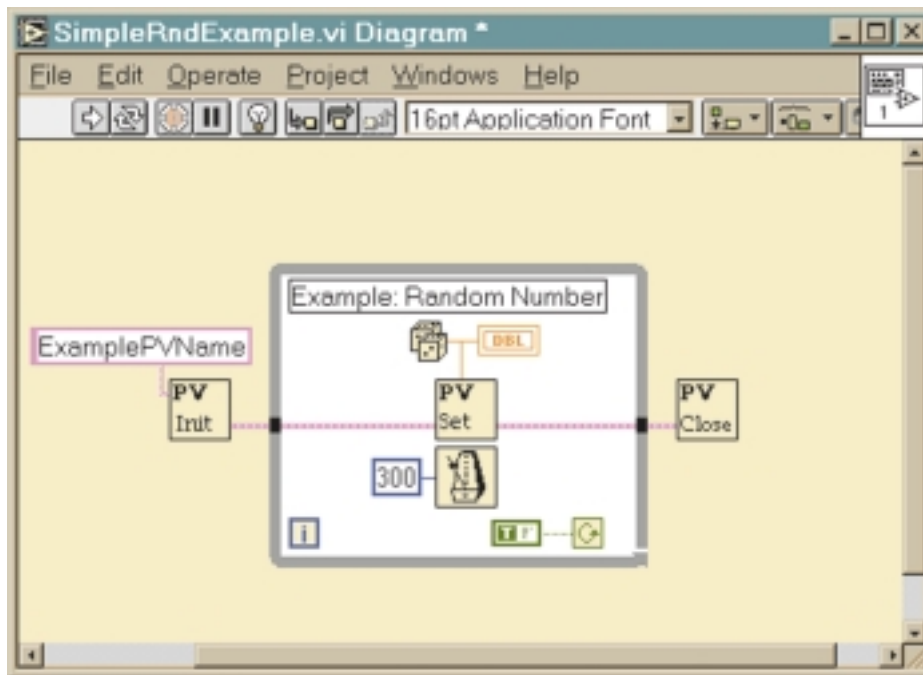
First experience shows that the performance and readability of LabVIEW code does not scale up, as well as vxWorks IOCs, which have successfully been configured to serve more than a thousand process variables. In LabVIEW, the programmer has full control over the implementation of scanning the hardware, responding to user input, and serving new values via CA, all in the familiar LabVIEW G schematic environment. For a more complex system, however, with many process variables at different update rates, operator input, save/restore mechanisms, and support for "bumpless" reboot, the LabVIEW programmer ends up re-implementing what the EPICS database handles transparently. On the other hand, little time is required to set up a small system in LabVIEW, and the resulting performance is often adequate for systems with a limited number of channels.

ChannelArchiver

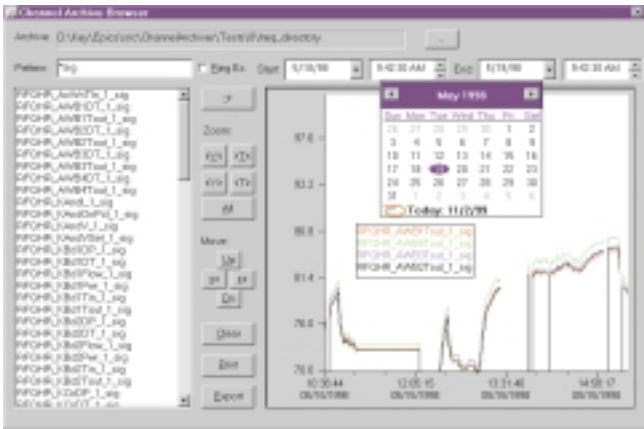
The CA is a generic, portable, networked archiving system that was developed for LEDA and has now become part of the EPICS toolkit.⁴ It stores a flat collection of channels, where a channel can be any process variable served by the CA. The archive engine can sample periodically at a fixed rate or on change. It saves the values together with a time stamp, status, alarm, and display information. A single-level hierarchy of channel groups can be archived conditionally, e.g., to stop archiving channels of a power supply when the device is off. The engine

has a built-in web server for access to status and configuration. It can store up to 10,000 values/s on a 450-MHz PC.

Several generic retrieval tools are available for data browsing: Native Unix and Win32 tools (Fig. 3), as well as a web interface (Fig. 4). The Unix tools were developed by Jefferson Laboratory and adapted to the CA. The Win32 and web interface were developed for LEDA. All these tools support plotting and can export data to spreadsheet programs. The native tools also allow interactive pan and zoom operations.



↑ Fig. 2. Minimal LabVIEW CA Server.



↑ Fig. 3. Archive data browser for Win32.

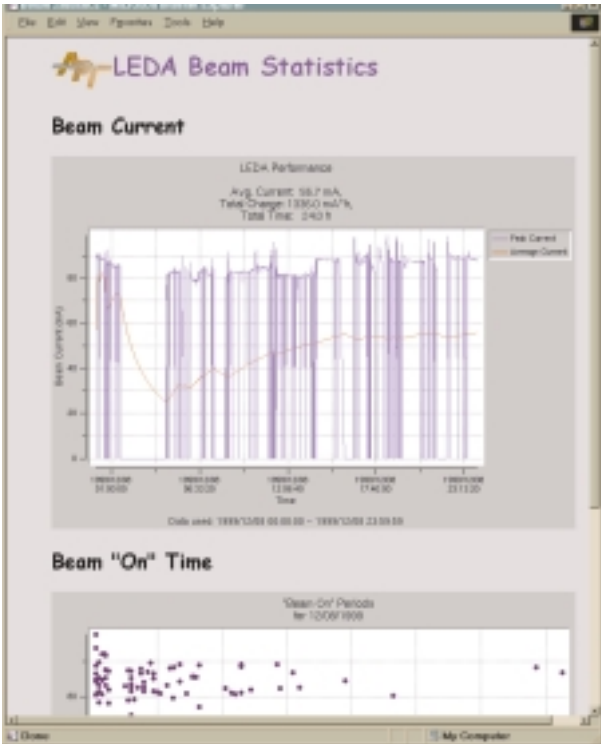


↑ Fig. 4. Web interface to archived data.

Because the engine stores the original time stamps as served by the IOCs with nanosecond granularity, values taken from multiple channels at the same nominal scan will rarely have exactly the same time stamps. Consequently, the retrieval tools provide generic filters for correlation: linear and staircase interpolation.

For more sophisticated data analysis, a portable scripting interface to the tool command language and python is used. Programmers can build on example scripts for text file export, time/Y plots, X/Y plots, etc. Customized scripts can therefore be written to answer questions such as: "How often was XX below 10.0 and for how long?"

Scripts can also generate daily statistics, including calculations for averages and integrals of the archived data, and publish them as web pages as shown in Fig. 5, a test setup for the automated generation of LEDA Beam Statistics. The scripting interface is available for Unix and Win32, as is the archive engine.



↑ Fig. 5. Web page created with script from archived data.

Conclusion

These new tools have been instrumental in integrating several LabVIEW-based components of the LEDA project into the distributed control system; they also helped to collect data for documenting the results.^{5,6,7} They will be used for future projects not only at LEDA⁸ but also at other EPICS sites. For example, the Swiss Light Source project at the Paul Scherrer Institute has been using the CA for approximately a year. The SNS project is planning to use the archiving system, and groups at the DESY and BESSY laboratories are trying to convert from their existing archiving system. With suggestions and contributions from these sites, both tools reflect the nature of EPICS as a living project that continues to grow with the requirements.

References

1. J. D. Schneider, "Operation of the Low-Energy Demonstration Accelerator: The Proton Injector for APT," in *Proceedings of the 1999 IEEE Particle Accelerator Conference* (IEEE Catalog No. 99CH36366, 1999), 503–507.
2. J. O. Hill, "EPICS Communication Loss Management," in *Proceedings of the 1993 International Conference on Accelerator and Large Experimental Physics Control Systems* (North Holland, 1994), 218–220.
3. J. O. Hill, "Channel Access: A Software Bus for the LAACS," in *Proceedings of the 1998 International Conference on Accelerator and Large Experimental Physics Control Systems* (North Holland, 1999), 352–355.
4. K. Kasemir and L. Dalesio, "Data Archiving in EPICS," in *1999 International Conference on Accelerator and Large Experimental Physics Control Systems*.
5. J. F. Power *et al.*, "Performance of the Beam Phase Measurement System for LEDA," Los Alamos National Laboratory document LA-UR-00-3367 (submitted to the Ninth Beam Instrumentation Workshop, May 8–11, 2000, Boston, Massachusetts).
6. D. Barr *et al.*, "LEDA Beam Diagnostics Instrumentation: Beam Current Measurement," (submitted to the Ninth Beam Instrumentation Workshop, May 8–11, 2000, Boston, Massachusetts).
7. D. Barr *et al.*, "LEDA Beam Diagnostics Instrumentation: Beam Position Monitors," (submitted to the Ninth Beam Instrumentation Workshop, May 8–11, 2000, Boston, Massachusetts).
8. L. A. Day *et al.*, "Control System for the LEDA 6.7-MeV Proton Beam Halo Experiment," Los Alamos National Laboratory document LA-UR-00-3842 (submitted to the Twentieth International Linac Conference, August 21–25, 2000, Monterey, California).

AccelView—An Innovative Tool Designed to Simulate, Visualize, and Animate Accelerator Beam Lines

C.K. Allen (LANSCE Division)

A software tool for the simulation, visualization, and animation of accelerator systems, known as AccelView, has been developed from support by the AHF project. Running on a Windows NT platform, AccelView comprises both an accelerator simulator and an animation system for real-time data, has a graphical user interface, and is event driven. Although still in the developmental stage, this new tool has already been used to analyze the imaging section of the proposed AHF accelerator. AccelView is extremely useful for demonstrating accelerator operation and in diagnosing unexpected beam-line behavior.

Computer Visualization and Animation


Computer visualization and animation have become important research tools in recent years—beam physics can find rich application for these technologies. Accelerator physicists must contend with enormous data sets; graphical representations of the data are a convenient way to cope with this fact. When data sets are also functions of time, one can animate the data to see correlations and trends in beam-line behavior that otherwise might be difficult to identify.

Visualization and animation are now widely accessible to both computer users and developers. For three-dimensional visualization and animation, the low-level graphics library OpenGL has become an industry standard.¹ This graphics library is designed to automatically recognize any specialized graphics hardware on the host platform. Originally developed by Silicon Graphics Incorporated (SGI), it is now available on most platforms, including MS Windows and Linux. There are other high-level, visualization libraries built over OpenGL, including Open Inventor² (by SGI) and the Visualization Toolkit.³

Fig. 1 shows one frame in a three-dimensional animation of a particle simulation using the OpenGL graphics library. The system being simulated is the periodic lattice of the LEDA beam-halo experiment.⁴ In the figure, two bunches can be seen travelling through the lattice, and one magnet surrounding the beam pipe is visible. The particles in the beam are color-mapped—red particles have large transverse velocities, blue particles have small transverse velocities, and white particles have intermediate velocities. Because there is no longitudinal focusing in the lattice, the beam is in the initial stages of debunching whereby the beam becomes continuous.

AccelView Application

AccelView, a computer system for the generalized simulation and animation of accelerator systems, is being developed through funds provided by the AHF project. This software tool has already been used to analyze the imaging section of the proposed AHF accelerator.⁵ AccelView runs on a Windows NT platform, has a graphical front end that supports multiple windows, and is event driven. By interacting with the user interface, one can manipulate various aspects of the simulation during run time, such as accelerator parameter



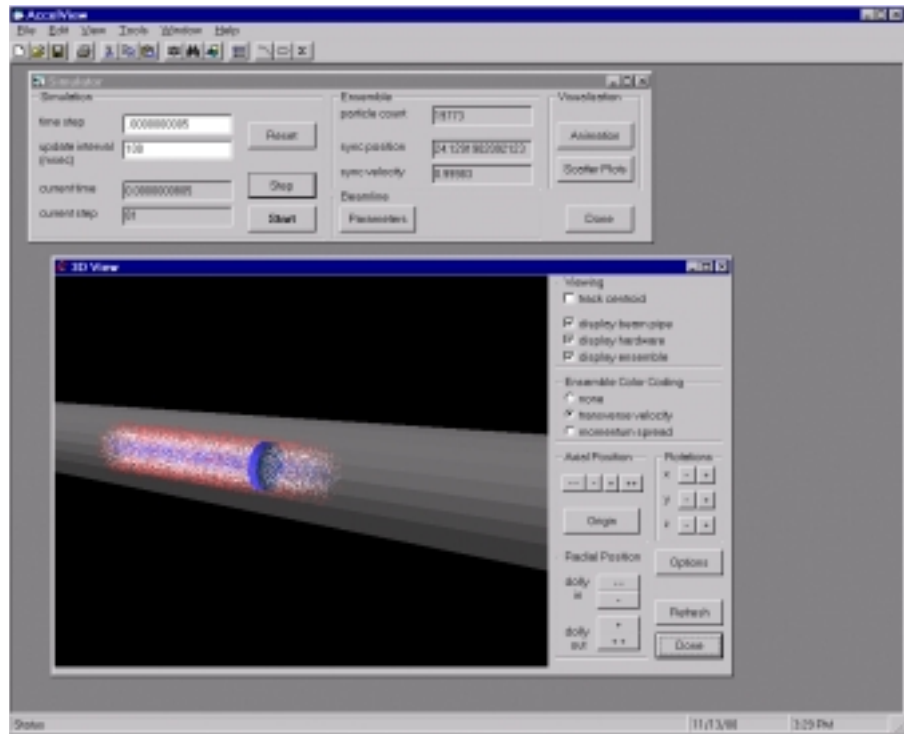
↑ Fig. 1. Visualization and animation of a particle-beam simulation using OpenGL.

values, particle distributions, update intervals, and other parameters. This situation is quite different from command-line simulators where input files are modified and simulations are run non-interactively.

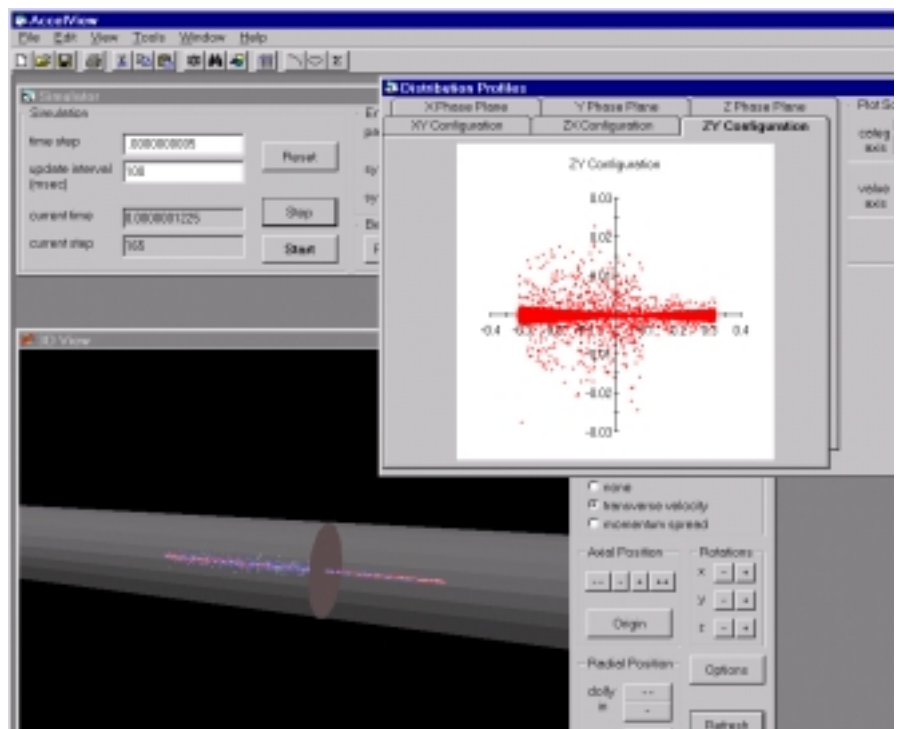
The particle simulation is controlled from the Simulator form, accessible from the Tools menu. From this form, the user can also launch different views of the simulation data, in particular, a three-dimensional rendering of the accelerator and particle ensemble. This view animates the current simulation data in real time—that is, we do not view a “movie” of simulation data that have been previously generated and stored. Moreover, the simulation may be stopped and restarted at any time to adjust the viewpoint, change parameters, or inspect different aspects of the simulation data (such as scatter plots or histograms).

Fig. 2 shows a snapshot of the AccelView application during a simulation. In the figure, we see two forms—the Simulator and the 3D View. In the Simulator form, we can change the time step and the update interval (for animation), as well as monitor simulation times and ensemble parameters. The 3D View form shows the beam bunch impinging upon a target; the cut-away beam pipe is also visible. As in Fig. 1, the particles are color-mapped according to transverse velocity. Also seen on the 3D View form are buttons for changing the camera position and several other viewing options.

The user is able to view different aspects of the simulation data at the same time. This feature is demonstrated in Fig. 3. At the



↑ Fig. 2. AccelView user screen shows a beam bunch striking a target.



↑ Fig. 3. AccelView user screen shows a beam bunch passing through a collimator.

instant depicted in the figure, the beam is passing through a collimator. Note in the figure, the additional form providing scatter plots of the ensemble in the various phase planes. The affect of the collimator is clearly seen in the YZ scatter plot.

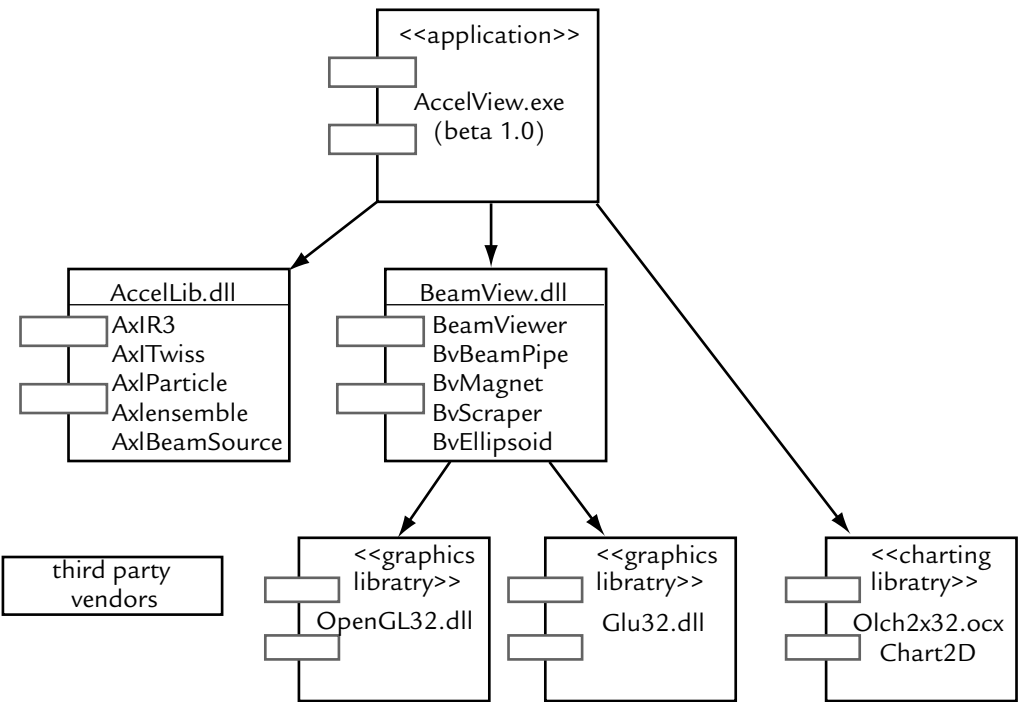
Simulation and Visualization Architecture

Charged-particle simulators typically are large software systems. Just as large hardware systems are designed *a priori* with blueprints and drawings (for example, accelerators), large software systems should be designed before they are implemented. This forethought becomes especially important if one is to add complex features such as visualization and animation of the simulation data. The Unified Modeling Language (UML)⁶ is a formal language for designing software systems and is the current standard for blueprinting software systems. A UML component diagram showing the basic structure of AccelView is presented in Fig. 4 where we see that the AccelView executable relies directly on the dynamic link libraries (DLLs), AccelLib.dll and BeamView.dll. The executable also directly employs Olch2x32.ocx, which is a third party library of ActiveX controls for two-dimensional charting. (There are several other DLLs, not shown, that AccelView relies upon for user interface functions and OLE Automation.) Also in the figure, we see that the

component BeamView.dll relies upon the dynamic-link libraries OpenGL32.dll and Glu32.dll.

Much of AccelView’s in-house implementation is contained in the two DLLs, AccelLib.dll and BeamView.dll. The two libraries contain software components for the simulation of accelerator systems and the visualization of such systems, respectively. The components are packaged using the Component Object Model (COM)⁷ and OLE Automation⁸ technologies of Microsoft for assembly in a rapid-application-development environment. In this type of design, precompiled software components are fit together in the aggregate system in much the same manner as integrated circuits are inserted into printed circuit boards. The convenience here is that, just as with integrated circuits, many third party software components may be used to assemble the complete application. Thus, only features specific to the particle simulation have to be implemented in house.

The three-dimensional visualization capability is contained in the BeamView library. The primary component of this library is an ActiveX control (called BeamViewer) that does the visualization. Internally, the control employs the OpenGL graphics library and its companion utility library GLU to perform the three-dimensional rendering. Also included in BeamView are many other software objects that represent hardware objects in the beam line, some of which are listed in Fig. 4. Each



↑ Fig. 4. AccelView, UML component diagram.

object knows how to render the hardware object it represents. Thus, displaying the accelerator hardware involves instantiating the corresponding software object and passing it to the ActiveX control for rendering.

Because of its component-based architecture, the implementation of AccelView was quite rapid. The majority of the effort for a single developer familiar with all the aforementioned software technology was about two man-months. In the development process, AccelView is currently in prototype form to demonstrate the visualization capabilities and the value of such a tool. New features are currently being added, mostly to enhance its generality and usefulness as a stand-alone interactive simulation tool.


References

1. M. Woo, J. Neider, and T. Davis, *OpenGL Programming Guide*, 2nd ed. (Addison-Wesley, Reading, Massachusetts, 1997).
2. J. Wernecke, *The Inventor Mentor: Programming Object-Oriented 3D Graphics with Open Inventor Release 2* (Addison-Wesley, Reading, Massachusetts, 1994).
3. W.J. Schroeder, K.M. Martin, and W.E. Lorensen, *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*, 2nd ed. (Prentice-Hall PTR, Upper Saddle River, New Jersey, 1998), Chap. 2.
4. T.P. Wangler, "LEDA Beam Halo Experiment—Physics and Concept of the Experiment," Los Alamos National Laboratory report LA-UR-00-3181 (2000).
5. P.W. Lisowski *et al.*, "AHF Technical Status Report," Los Alamos National Laboratory report LA-UR-01-838 (2001).
6. G. Booch, J. Rumbaugh, and I. Jacobson, *The Unified Modeling Language User Guide* (Addison-Wesley, Reading, Massachusetts, 1999).
7. D. Rogerson, *Inside COM: Microsoft's Component Object Model* (Microsoft Press, Redmond, Washington, 1997).
8. K. Brockschmidt, *Inside OLE*, 2nd ed. (Microsoft Press, Redmond, Washington, 1995).

H⁻ Ion Source Development at LANSCE

W.B. Ingalls, R.L. Haffner, E.G. Jacobson, B. J. Meyer, B.A. Prichard, O.R. Sander, J.D. Sherman, J.E. Stelzer, R.R. Stevens, J.D. Weiting (LANSCE Division), K.N. Leung, R.W. Thomae, M.D. Williams (Lawrence Berkeley National Laboratory)


A new cusp-field, surface-conversion ion source for the production of H⁻ beam suitable for injection into the LANSCE 800-MeV linac has been developed. The new ion source is the result of a collaboration between personnel at LBNL and LANL personnel in LANSCE-2. The collaboration is a natural one; the existing H⁻ ion source at LANSCE evolved from an earlier ion source developed at LBNL. The new ion source is capable of producing H⁻ beam currents of 40 mA, more than twice the beam intensity routinely produced by the present production ion source. The increased intensity of the beam available for acceleration to 800 MeV by the LANSCE linac is necessary to achieve the Short-Pulse Spallation Source (SPSS) project goal of 200 μ A in the PSR. Fig. 1 is a photograph of the new ion source. The new ion source is referred to as the Y2K source to distinguish it from the present production ion source (shown in Fig. 2).



↑ Fig. 1. The Y2K ion source. Six filaments can be seen in the source interior. The back plate with a converter mounted in it is seen to the lower right. The ion beam is formed and extracted along the axis of the source body.

New Ion Source Development

The two ion sources are, in principle, very similar. Both employ an electrode referred to as the converter to form the H⁻ beam. Cesium metal is introduced into both ion sources to coat the spherically shaped surface of the converter where H⁺ ions from the ion source plasma are converted into H⁻. The plasma is confined within the fields generated by the cusp magnets



↑ Fig. 2. The LANSCE production ion source. A filament side plate has been removed to show the source interior. The beam in this source exits longitudinally through the conically shaped aperture in the side of the ion source body.

arrayed around the outside of the source body in both ion sources. These magnets are easier to see in the production source photograph. The ion sources differ in the number and arrangement of the filaments that supply electrons to generate the ion source plasma and in the way the generated H⁻ beam leaves the ion source. The Y2K source has six filaments arrayed around the axis of the plasma chamber. In the LANSCE production source, there are two filaments mounted at either end of the source chamber. In the production source, the ion beam leaves the source through a hole in the side of the ion source chamber, whereas the ion beam leaves the Y2K source on the source chamber axis.

One of the primary constraints placed on the new ion source was the requirement that it be possible to install it in the equipment dome of the existing H⁻ Cockcroft-Walton injector so that it could be mounted as a direct replacement for the present production ion source. The higher beam current available from the new ion source precipitated the redesign of the 80-kV accelerating column in the equipment dome and an upgrade to the electronics needed to operate the ion source. LBNL personnel performed the physics design of the new column; the engineering details were done within LANSCE-2. The new column electrode structure fits in the existing column ceramic insulator, resulting in an accelerating column that is outwardly no different from


the column presently in use. The new electronics have been designed and constructed by personnel in LANSCE-6.

To ensure there are relatively few surprises when the new ion source and column are put into service at LANSCE, we have constructed a test stand to evaluate them offline. Early on, it was recognized that the ISTS should duplicate the existing beam transport equipment in the H^- injector as closely as possible to minimize the chance that the new equipment will not perform as expected when installed. ISTS has been built in the location of an older test stand by the same name. The 2000 l/s turbo-molecular vacuum pump and the attendant Roots blower and forepump from the old stand were used on the new ISTS. A solenoid lens from the old test stand was modified to duplicate the first of the two solenoid lenses found in the H^- injector dome transport.

The remainder of the ISTS beam transport has been assembled largely from H^- injector spare parts and other parts fabricated to duplicate those in the H^- injector. The ISTS incorporates beam diagnostics in the form of two additional emittance stations not found in the H^- injector. A rendering of the layout of ISTS is shown in Fig. 3. There is an emittance station to measure the beam phase-space immediately after the 80-kV accelerating column, a second station that duplicates the dome emittance station at the midpoint of the transport, and a third station for measurements after the bending magnet and second beam transport solenoid lens.

Some of the early testing of the new 80-kV column on ISTS revealed an electron-heating problem. The lens elements that warped due to this heating were redesigned, and the new column is now performing as expected. We have discovered a similar overheating problem with the Y2K ion source that results in undesirable movement and shorting of internal components. While a solution to the cooling problem is being developed, other aspects of the ion source that affect performance have been under investigation with a proof-of-principle ion source capable of producing as much as 35 mA of H^- beam.

We have learned that the configuration of the magnetic fields in the repeller electrode internal to the ion source can have a dramatic effect on the amount of electron suppression by the repeller. A cusp-field configuration, as is used in the production ion source, results in an electron-to-beam ratio of 4.4:1. This ratio dropped to 2.3:1 when the configuration of the magnets was changed to a solenoidal field. A further reduction to



↑ **Fig. 3.** Artist's rendering of ISTS. The beam transport progresses from left to right. In this view, the ion source is located at the left of the figure but is largely hidden by the 80-kV column insulating jacket. The emittance gear actuators protruding from the first, second, and third beam diagnostic stations are clearly evident. The first station actuators, being located on ports in other apparatus, are in the horizontal plane only.

1:1 was obtained with an undulator configuration of magnets. Both the solenoid and undulator configuration resulted in reductions of 10% to 20% in the beam emittance, a measure of the beam quality, compared to the cusp-magnet configuration.

The emittance contour-plots obtained on ISTS sometimes indicate the presence of two beams, both H^- , but formed in different ways in the ion source. The less-intense second beam had been observed in the past but had been attributed to the formation of a neutral beam by charge-exchange of the primary beam with the residual background gas in the ion source and beam transport. On ISTS, the second beam is observed at the third emittance station, after the beam has passed through a magnetic field. This observation indicates the beam is charged, not neutral. The appearance of the second beam has been simulated with the ion beam code PBGUNS. The code has been used to demonstrate minor modifications to the beam extraction region of the ion source to mitigate the formation of the second beam.

Because the ISTS consists of an ion source mounted on a vacuum system and beam transport that is virtually identical to that found in the LANSCE H^- injector, it is proving to be an ideal locale to test the new electronics that are also a part of the SPSS Enhancement Project (see related story on page 142).

Chopper Control

R.O. Nelson, R.B. Merl, C.R. Rose (LANSCE Division)

Through careful target and moderator design at a spallation source, we tailor the neutron energy spectrum to match the requirements of the neutron-scattering instruments. But there are limits to what we can accomplish with these techniques. We often want to further filter the neutron beam as it emerges from the neutron source and begins its journey through the instrument. Neutron choppers are electro-mechanical devices that can remove undesired energy components from the beam. So-called time zero or T-zero (T_0) choppers rotate a large mass through the beam to effectively place a beam stop in the path of the beam to eliminate the high-energy neutrons that occur early in the neutron pulse. We use another type of chopper to block most neutrons except those in a narrow low-energy band. Generally, we place this energy selector (E_0) chopper downstream from a companion T_0 chopper. We must operate such choppers in phase with the production of the neutron pulse and in phase with one another so that the energy distribution in each neutron pulse in the instrument remains constant.


Unfortunately, accelerator power systems have not always allowed us to deliver the proton beam to the target in a periodic manner. Fluctuations in the power grid give rise to fluctuations in the time of arrival of the protons on target and hence to the time of arrival of the neutrons in the instrument. To maintain a constant neutron energy spectrum, we must continually alter the phase of the rotating chopper blades to match the phase of the accelerator, which follows the power grid. Because the choppers consist of the massive blades rotating in the neutron beam and lethargic motors, they cannot easily accommodate these phase changes. On the other hand, the power modulators in the accelerator operate best when they are able to follow phase fluctuations in the power grid. Clearly, a synchronization problem exists if the accelerator phases beam delivery with the power grid while the choppers operate at a fixed frequency. Hence over the past 20 years, we have designed accelerator and chopper controls that are engineering compromises that work adequately for neutron-scattering facilities like LANSCE¹⁻³ and ANL.⁴⁻⁶

During an upgrade to the LANSCE systems and for the design of the timing system for the new SNS, we have reviewed the engineering compromises and, with better


analysis tools, have found new options based on new technologies that are available for high-performance implementation. We present our system-engineering analysis and experimental results obtained with prototype systems.

Predicting and Analyzing Chopper Performance

To predict and analyze neutron-chopper performance, we modeled all subsystems beginning at the power grid and ending at the rotating chopper blade. All simulations were performed with MATLAB and SIMULINK software packages running on PCs. The model has four major components: the power grid, the timing-reference generator, the chopper controllers, and the choppers (Fig. 1). Initially, we measured phase variations in the power grid using a zero-crossing detector and a global-positioning-system, time-stamping technique. We then developed a digital-signal-processing (DSP) algorithm (based on three-phase, power-line measurements) that extracts the grid phase drift measured with respect to an ideal 60-Hz reference.⁷ Using the DSP algorithm, we recorded phase-drift data that we subsequently used as input into our simulations. Fig. 2 shows the phase-drift data recorded during a 90-h period; these data dramatically make the point that the phase of the power grid wanders. Although it is easy to work with, it does bandwidth limit the input data because of its weighted least-squares algorithm. The proposed timing-reference generator takes the phase information from the grid and smoothes it. We vary the smoothing parameters to effectively limit the




↑ Fig. 1. Block diagram for timing-system dependencies. Accelerator and chopper controllers act strictly as slaves to the master timing-reference generator.



↑ **Fig. 2.** Ninety hours of phase-drift data. The drift is measured with respect to an ideal 60-Hz source. Note that the drift is substantial (1500 radians corresponds to 240 cycles or 4 s on a motor-driven clock). The challenge for the chopper is to follow these changes through acceleration and deceleration.

phase acceleration passed along to the choppers. Fig. 3 illustrates the concept of smoothing, starting with the waveform from the power grid. The zero-crossing pulses mark the phase of the grid as it passes through half cycles. We have developed several algorithms to generate the smoothed timing-reference pulses. Each algorithm forms a phase difference by subtracting the smoothed phase from the raw grid phase. The phase difference is then processed with various filters and proportional-integral compensators. This combination produces a correction signal that modulates the frequency of the software-controlled frequency generator, which has a center frequency of 60 Hz.

We used 20-Hz T_0 , 60-Hz T_0 , and 600-Hz E_0 choppers (each paired with its own controller) to realistically model a neutron-scattering facility. Because controllers for frame-overlap choppers are identical to T_0 -chopper



↑ **Fig. 3.** Illustration of key timing concepts. The timing-reference signal is a “smoothed” version of the zero-crossing signal that is appropriate for controlling choppers and simultaneously satisfies the accelerator.

controllers except for the tuning required for the smaller moments of inertia, frame-overlap choppers were omitted from the modeled systems. To compare the results of simulations with actual hardware performance, we constructed prototype controllers for T_0 choppers and operated them with chopper hardware already installed at the Lujan Center. E_0 -chopper data came from the recently purchased system from Revolve Magnetic Systems, Inc.

Because the performance of the various choppers, as well as the performance of the accelerator, is linked to the coupling to the grid, we quantified this coupling by measuring the standard deviation for the timing-reference phase with respect to the raw grid phase. We defined this metric as the accelerator-tracking sigma and varied it by changing the smoothing parameters in the timing-reference generator. Traditionally, accelerators prefer a small accelerator-tracking sigma (tight grid coupling), whereas choppers prefer a large accelerator-tracking sigma (loose grid coupling). In addition, to quantify chopper performance, we measured the standard deviation for the chopper arrival at top dead center with respect to the timing-reference pulse. Analogously, we defined this metric as the chopper-tracking sigma. Fig. 3 illustrates both the accelerator- and chopper-tracking intervals.

As a function of increasing accelerator-tracking sigma, we expect the chopper-tracking sigma to decrease. Increasing accelerator-tracking sigma allows the timing-reference phase to wander further from the grid phase, thereby averaging out some of the high-frequency variations of the grid. Thus the choppers have an easier time following only the slower variations of the grid. In the limit for large accelerator-tracking sigma, the timing-reference generator ignores the grid, and the choppers run at a constant timing reference frequency, 60 Hz.

Our goal was to calculate and then measure the chopper-tracking sigmas as a function of the accelerator-tracking sigma. Accelerator designers could then immediately determine the consequence on chopper performance for any choice of accelerator-tracking sigma. The chopper veto rate was derived from the chopper-tracking sigma when the veto window was specified and the distribution was Gaussian. Preliminary data suggested a nearly Gaussian distribution.

Simulations and Measurements

The simulations began with the choppers at rest, and it took approximately 300 s for them to come up to

speed. The simulation ended at 700 s, and we selected the data to be representative of typical power-grid behavior. The simulations were performed as a two-step process. Using the grid-phase drift data captured earlier, the first step smoothed data subject to the selected parameters, calculated the accelerator-tracking sigma for the smoothing process, and wrote out a test-vector file for use in the second step. There the file was used as the timing reference for all the choppers in the test suite. Each chopper and controller responded to the identical stimulus so that comparisons between choppers were valid.

During the simulations, the chopper-tracking sigma was recorded for each chopper- and test-vector combination, and the results were summarized (Fig. 4). One key aspect of the high performance of the high-speed chopper is that its rotor moment of inertia is about $0.01 \text{ kg}\cdot\text{m}^2$, whereas the moments of inertia of the low-speed choppers range from 8 to $12 \text{ kg}\cdot\text{m}^2$. The smaller moment of inertia of the high-speed chopper allows for much greater chopper bandwidth and thus lower timing jitter.

Discussion

The data from the simulations shown in Fig. 4 produce the expected trends. If the timing reference generator smooths out variations in the grid, then the choppers can follow the smoothed phase drift better. Greater smoothing produces still smaller chopper-tracking sigmas. However, as the timing-reference generator increases the smoothing, the chopper-tracking sigmas asymptotically approach a value near zero. They will not approach zero because of the system noise, quantizing errors, and control-circuit internal timing jitter. The simulations for the 600-Hz chopper produce


sigmas that are far too optimistic ($< 100 \text{ ns}$) in comparison with measured data (600 ns at an accelerator-tracking sigma of $22 \mu\text{s}$). Because the model ignores sampling errors, truncation of precision, and noise, it is not surprising that such a small sigma would be invalid. However, this modeling problem is not a serious concern overall because the measured 600-Hz chopper-tracking sigmas make no significant contribution to the resolution of the instruments.

The simulations show that the phase errors for all choppers are similar in magnitude and are coherent. To understand this effect, consider the case where all the choppers are synchronized to the grid. Now, suppose the phase of the grid randomly shifts. Initially, all the choppers are left with the same phase error. When they catch up as best they can, the grid moves again and the process repeats itself. Because the phase errors are similar, the chopper-tracking sigmas for choppers that operate at different velocities should scale in inverse proportion to their velocities. For example, the chopper-tracking sigma of a chopper operating at 60 Hz should be one-third the chopper-tracking sigma of a chopper operating at 20 Hz. This trend is seen in Fig. 4.

Because grid variations uniformly alter the phase error of choppers running at different speeds and non-uniformly alter the timing errors, techniques for extracting beam from the PSR or for triggering the accelerator may improve the timing for choppers at one speed while degrading the timing at other speeds. Because the chopper-tracking sigmas are smallest for high-speed choppers, such techniques can improve the already high performance of these choppers without adversely impacting the slower choppers. The inverse is not true. Therefore, these techniques are of marginal value unless extreme precision is required. Generally, other timing problems (e.g., moderator holdup times) would render this precision meaningless.

Conclusions

Low-speed choppers present the worst timing-synchronization problems because of their high moment-of-inertia rotors. Often, the question is posed whether or not additional motor torque would help. We have verified through models and empirical data that our motors have ample torque. In other words, more torque will not reduce the timing jitter. The fundamental problem is control-system bandwidth as it relates to moment-of-inertia rotor and to the loop stability at high gain.



↑ Fig. 4. The results of the simulations and measurements. The lines show predictions for the chopper suite. At this time, we have data for only the 20-Hz and 600-Hz choppers.

For LANSCE, the current timing-reference controller³ that gives an accelerator-tracking sigma of 25 μ s is good enough. Our chopper controllers have improved so much that obtaining another factor of two in the chopper-tracking sigma for the 20-Hz chopper would require the accelerator-tracking sigma to increase by a factor of four. Accelerator operations will not permit such a change. Furthermore, we chose to continue operations with an accelerator-tracking sigma of 25 μ s because the chopper-tracking sigmas are well matched to other resolution-limiting characteristics of instruments and moderators.

For SNS, an accelerator-tracking sigma of 125 μ s is good. Chopper-tracking sigmas are well matched to other resolution-limiting characteristics. An accelerator-tracking sigma of 125 μ s allows a 20-Hz chopper

with a large moment of inertia to operate with a chopper-tracking sigma of 7 μ s. From the accelerator perspective, operating with this accelerator-tracking sigma allows the timing-reference generator to be within ± 4 sigma or ± 500 μ s of the power grid 99% of the time. Increasing the accelerator-tracking sigma by a factor of four only improves the chopper-tracking sigma by a factor of two.

For LANSCE and SNS, the choppers and the accelerator should be phased to a master timing-reference generator, and the complications of ring extraction and beam triggering should not be dealt with. However, the accelerator must deliver the beam to the target at a fixed phase.

References

1. V. W. Bolie, R. M. Brugger, and R. N. Silver, "Neutron Chopper Phasing for a Quasi-Periodic Pulsed Neutron Source," *Nuclear Instruments and Methods in Physics Research A* **236**:1, 85-94 (1985).
2. M. Nutter, L. Lewis, S. Tepper, R. N. Silver, and R. H. Heffner, "Neutron Chopper Development at LANSCE," in *Proceedings of the Eighth Meeting of the International Collaboration on Advanced Neutron Sources* (Oxford, United Kingdom, July 8-12, 1985), RAL-85-110.
3. L. J. Rybarczyk and F. E. Shelley, Jr., "Improvements to the LANSCE Accelerator Timing System," 1997 Particle Accelerator Conference (Vancouver, British Columbia, May 12-16, 1997).
4. A. Rauchas, G. Ostrowski, C. Pelizzari, and G. Volk, "IPNS Accelerator System and Neutron Chopper Synchronization," in *Proceedings of the Seventh Meeting of the International Collaboration on Advanced Neutron Sources* (Chalk River, Ontario, Canada, September 13-16, 1983), AECL-8488.
5. L. I. Donley, "Phase Locking the IPBS Neutron Choppers to the 60-Hz Power Line," in *Proceedings of the Eighth Meeting of the International Collaboration on Advanced Neutron Sources* (Oxford, United Kingdom, July 8-12, 1985), RAL-85-110.
6. G. E. Ostrowski, L. I. Donley, A. V. Rauchas, G. J. Volk, E. A. Jung, J. R. Haumann, and C. A. Pelizzari, "The PNS Chopper Control and Accelerator Interface Systems," in *Proceedings of the Eighth Meeting of the International Collaboration on Advanced Neutron Sources* (Oxford, United Kingdom, July 8-12, 1985), RAL-85-110.
7. H.-S. Song, H.-G. Park, and K. Nam, "An Instantaneous Phase Angle Detection Algorithm Under Unbalanced Line Voltage Condition," in *Thirtieth Annual IEEE Power Electronics Specialists Conference* (Charleston, South Carolina, June 27-July 1, 1999).

